

A Flexible Krylov Solver for Shifted Systems with Application to Oscillatory Hydraulic Tomography

Arvind K. Saibaba* Tania Bakhos* Peter K. Kitanidis^{†*}

December 18, 2012

Abstract

We discuss efficient solutions to systems of shifted linear systems arising in computations for oscillatory hydraulic tomography (OHT). The reconstruction of hydrogeological parameters such as hydraulic conductivity and specific storage, using limited discrete measurements of pressure (head) obtained from sequential oscillatory pumping tests, leads to a nonlinear inverse problem. We tackle this using the quasi-linear geostatistical approach [13]. This method requires repeated solution of the forward (and adjoint) problem for multiple frequencies, for which we derive flexible preconditioned Krylov subspace solvers specifically for shifted systems. The solvers allow the preconditioner to change at each iteration. We analyse the convergence of the solver and perform an error analysis when an iterative solver is used for inverting the preconditioner matrices. Finally, we apply our algorithm to a challenging application taken from oscillatory hydraulic tomography to demonstrate the computational gains by using the resulting method.

1 Introduction

Hydraulic tomography (HT) is a method for characterizing the subsurface that consists of applying pumping in wells while aquifer pressure (head) responses are measured. Using the data collected at various locations, important aquifer parameters (e.g., hydraulic conductivity and specific storage) are

*Institute for Computational and Mathematical Engineering, Stanford University

[†]Department of Civil and Environmental Engineering, Stanford University

estimated. An example of such a technique is transient hydraulic tomography (reviewed in [3]). With transient hydraulic tomography the measured signals after a certain distance is typically weak in comparison to ambient noise (i.e., disturbances from pumping other other wells, changes in water level in adjacent streams, etc.). Oscillatory hydraulic tomography (OHT) is an emerging technology for aquifer characterization that involves a tomographic analysis of oscillatory signals. Here we consider that a sinusoidal signal of known frequency is imposed at an injection point and the resulting change in pressure is measured at receiver wells. Consequently, these measurements are processed using a nonlinear inversion algorithm to recover estimates for the desired aquifer parameters. Oscillatory hydraulic tomography has notable advantages over transient hydraulic tomography; namely, a weak signal can be distinguished from the ambient noise and by using signals of different frequencies, we are able to extract additional information without having to drill additional wells.

Using multiple frequencies for OHT has the potential to improve the quality of the image. However, it involves considerable computational burden. Solving the inverse problem, i.e. reconstructing the hydraulic conductivity field from pressure measurements, requires several application of the forward (and adjoint) problem for multiple frequencies. As we shall show in section 5, solving the forward (and adjoint) problem involves the solution of shifted systems for multiple frequencies. For finely discretized grids, the cost of solving the system of equations corresponding to each frequency can be high to the extent that it might prove to be computationally prohibitive when many frequencies are used, for example, on the order of 200. The objective is to develop an approach in which the cost of solving the forward (and adjoint) problem for multiple frequencies is not significantly higher than the cost of solving the system of equations for a single frequency - in other words, the cost should depend only weakly on the number of frequencies.

Direct methods, such as sparse LU, Cholesky or LDL^T factorization, are suited to linear systems in which the matrix bandwidth is small, so that the fill-in is somewhat limited. An additional difficulty that direct methods pose is that for solving a sequence of shifted systems, the matrix has to be re-factorized for each frequency, resulting in a considerable computational cost. By contrast, Krylov subspace methods for shifted systems are particularly appealing since they exploit the shift-invariant property of Krylov subspaces [28] to obtain approximate solutions for all frequencies by generating a single approximation space that is shift independent. Several algorithms have been developed for dealing with shifted systems. Some are based on Lanczos recurrences for symmetric systems [19, 20]; others use the unsymmetric Lanczos [9], and some others use Arnoldi iteration [6, 10, 26, 5, 12].

Shifted systems also occur in several other applications such as control theory, time dependent partial differential equations, structural dynamics, and quantum chromodynamics (see [26] and references therein). Hence, several other communities can benefit from advances in efficient solvers for shifted systems. The Krylov subspace method that we propose is closest in spirit to [12]. However, as we shall demonstrate, we have extended their solver significantly.

Contributions: Our major contributions can be summarized as follows:

- We have extended the flexible Arnoldi algorithm discussed in [12] for shifted systems of the form $(K + \sigma_j M)x_j = b$ for $j = 1, \dots, n_f$ that employs multiple preconditioners of the form $(K + \tau M)$. Further, analyze the convergence of the solver.
- When an iterative solver is used to apply the preconditioner, we derive an error analysis that gives us stopping tolerances for monitoring convergence without constructing the full residual.
- Our motivation for the need for fast solvers for shifted systems comes from oscillatory hydraulic tomography. We describe the key steps involved in inversion for oscillatory hydraulic tomography, and discuss how the computation of the Jacobian can be accelerated by the use of the aforementioned fast solvers.

Limitations: The focus of this work has been on the computational aspects of oscillatory hydraulic tomography. Although the initial results are promising, several issues remain to be resolved for application to realistic problems of oscillatory hydraulic tomography. For example, we are inverting for the hydraulic conductivity assuming that the storage field is known. In practice, the storage is also unknown and needs to be estimated from the data as well. Moreover, simulating realistic conditions (higher variance in the log conductivity field, and adding measurement noise in a realistic manner) may significantly improve the error with the addition of information from different frequencies. We will deal with these issues in another paper.

The paper is organized as follows. In section 2, we discuss the Krylov subspace methods for solving shifted linear systems of equations based on the Arnoldi iteration using preconditioners that are also shifted systems. In section 4, we discuss an error analysis when an iterative method is used to invert the preconditioner matrices. In section 5, we discuss the basic constitutive equations in OHT, which can be expressed as shifted linear system of equations and discuss the geostatistical method for solving inverse problems. Finally, in section 6 we present some numerical results on systems of shifted

systems and then discuss numerical results involving the inverse problem arising from OHT. We observe significant speed-ups using our Krylov subspace solver.

2 Krylov subspace methods for shifted systems

The goal is to solve systems of equations of the form

$$(K + \sigma_j M) x_j = b \quad j = 1, \dots, n_f \quad (1)$$

Note that σ_j , for $j = 1, \dots, n_f$ are (in general) complex shifts. We assume that none of these systems are singular. In particular, for our application, both K and M are stiffness and mass matrices respectively and are positive definite but our algorithm only requires that they are invertible. By using a finite volume or lumped mass approach, the mass matrices become diagonal but this assumption is not necessary. Later, in sections 5.1 and 5.3, we will show how such equations arise in our applications.

Several efficient methods exist for solving the system of equations (1), which solve for multiple shifts roughly at the cost of a single system. They do so by generating a subspace that is independent of the shift and use the shift-invariant property of Krylov subspaces (for a detailed review, see [28, section 14.1] and references therein). However, in practice, the number of iterations taken can often be large, especially for large matrices from realistic applications.

In order to minimize the number of iterations, Meerbergen [19] proposes a left preconditioner of the form $K_\tau \stackrel{\text{def}}{=} K + \tau M$ that is factorized and inverted using a direct solver. The application of K_τ^{-1} to a vector is, in general, not cheap but the spectrum of $(K + \tau M)^{-1}(K + \sigma M)$ is often more favorable, which results in fast convergence of the Krylov methods in just a few iterations [19]. This form of preconditioning has its roots in solving large-scale generalized eigenvalue problems and is known as *Cayley transformation* [11]. In [21], the authors provide some analysis for choosing the best value of τ that optimally preconditions all the systems. However, we observed that (also, see [12]) using a single preconditioner for all the systems may not yield optimal convergence for all systems. In [12], the authors propose a flexible Arnoldi method for shifted systems that uses different values of τ resulting in different preconditioners at each iteration. This can potentially reduce the number of iterations for all the shifted systems. Before we describe our flexible algorithm in section 2.2, we will derive the right preconditioned version

of Krylov subspace method for shifted systems. This serves two purposes - it motivates our algorithm, while clarifying some of the notation.

2.1 Right preconditioning for shifted systems

As mentioned earlier, we will review the right preconditioned version of the Krylov subspace algorithm for shifted systems. Following the approach in [19, 28], we solve the system of equations (1) using a shifted right preconditioner of the form $K_\tau \stackrel{\text{def}}{=} K + \tau M$

$$(K + \sigma_j M) K_\tau^{-1} \bar{x}(\sigma_j) = b \quad x(\sigma_j) = K_\tau^{-1} \bar{x}(\sigma_j) \quad (2)$$

for $j = 1, \dots, n_f$.

Algorithm 1 Arnoldi using Modified Gram-Schmidt [24]:

- 1: Given M , b the right hand side.
 - 2: Compute $v_1 = b/\beta$ and $\beta \stackrel{\text{def}}{=} \|b\|_2$
 - 3: Choose τ and factorize $K_\tau \stackrel{\text{def}}{=} K + \tau M$
 - 4: Define the $m+1 \times m$ matrix $\bar{H}_m = \{h_{i,k}\}_{1 \leq i \leq m+1, 1 \leq k \leq m}$. Set $\bar{H}_m = 0$.
 - 5: **for all** $k = 1, \dots, m$ **do**
 - 6: Compute $z_k = K_\tau^{-1} v_k$
 - 7: $w_k := M z_k$
 - 8: **for all** $i = 1, \dots, k$ **do**
 - 9: $h_{ik} := w_k^* v_i$
 - 10: Compute $w_k := w_k - h_{ik} v_k$
 - 11: **end for**
 - 12: $h_{k+1,k} := \|w_k\|_2$. If $h_{k+1,k} = 0$ stop
 - 13: $v_{k+1} = w_k / h_{k+1,k}$
 - 14: **end for**
 - 15: Define $V_{m+1} = [v_1, \dots, v_{m+1}]$ and $Z_m = [z_1, \dots, z_m]$
-

The algorithm proceeds as follows: first, we run m steps of the Arnoldi algorithm on the matrix $M K_\tau^{-1}$. This is summarized in algorithm 1. At the end of m steps or the Arnoldi process, we have the following relations

$$M Z_m = V_{m+1} \bar{H}_m \quad (3)$$

$$(K + \tau M) Z_m = V_m \quad (4)$$

where, $V_m^T V_m = I$. Multiplying the first equation by $(\sigma_j - \tau)$ and adding it to the second equation gives us

$$(K + \sigma_j M)Z_m = V_{m+1} \underbrace{\left(\begin{bmatrix} I \\ 0 \end{bmatrix} + (\sigma - \tau)\bar{H}_m \right)}_{\stackrel{\text{def}}{=} \bar{H}_m(\sigma_j; \tau)} = V_{m+1} \bar{H}_m(\sigma_j; \tau) \quad (5)$$

We seek solutions of the form $x_m = Z_m y_m$ (with zero as the initial guess), since by construction, Z_m forms a basis for the Krylov subspace $\mathcal{K}_m(MK_\tau^{-1}, b)$, where

$$\mathcal{K}_m(A, b) = \text{span}\{b, Ab, \dots, A^{m-1}b\}$$

By minimizing the residual norm over all possible vectors in $\mathcal{K}_m(MK_\tau^{-1}, b)$, we obtain the generalized minimum residual (GMRES) method for shifted systems, whereas by imposing the Galerkin condition $r_m \perp \mathcal{K}_m(MK_\tau^{-1}, b)$, we obtain the full orthogonalized method (FOM) for shifted systems. This is summarized in algorithm 2.

Algorithm 2 FOM/GMRES for Shifted Systems

- 1: Given matrices K and M , a right hand side b , $\sigma \in \{\sigma_1, \dots, \sigma_{n_f}\}$
 - 2: Choose τ , build $K_\tau \stackrel{\text{def}}{=} K + \tau M$ and construct preconditioner.
 - 3: **for all** $l = 1, 2, \dots, m$ **do**
 - 4: Run m steps of algorithm 1 to generate V_{m+1}, \bar{H}_m and Z_m .
 - 5: **for all** $j = 1, \dots, n_f$ **do**
 - 6: If, not converged
 - 7: Construct $\bar{H}_m(\sigma_j; \tau)$ as defined in equation (5).
 - 8: FOM:

$$H_m(\sigma_j; \tau)y_m^{fom}(\sigma_j) = \beta e_1$$
 - 9: GMRES:

$$y_m^{gm}(\sigma_j) \stackrel{\text{def}}{=} \min_{y_m \in \mathbb{C}^m} \|\beta e_1 - \bar{H}_m(\sigma_j; \tau)y_m\|_2$$
 - 10: Construct the approximate solution as $x_m(\sigma_j) = Z_m y_m(\sigma_j)$
 - 11: **end for**
 - 12: **end for**
-

Thus, there is a distinct advantage in using iterative solvers for shifted systems; the expensive step of constructing the basis for the Krylov subspace is performed only once and using the shift-invariance property of the Krylov subspace, the sub-problem for each shift in 2 can be computed at relatively low cost. In [19], the spectrum of $(K + \sigma M)(K + \tau M)^{-1}$ was analyzed and it was shown that the preconditioner K_τ is well suited only for values of

frequencies σ near τ . However, the values of σ can be widely spread and a single preconditioner K_τ might not be a good choice for preconditioning all the systems. In section 6.1, we demonstrate an example in which a single preconditioner does not satisfactorily precondition all the systems. In [12], the authors propose a flexible approach using a (possibly) different preconditioner at each iteration. We shall adopt this approach.

2.2 Flexible preconditioning

We now describe our flexible Krylov approach for solving shifted systems. Following Saad [23] and [12], we use a variant of GMRES which allows a change in the preconditioner at each iteration. In algorithm 1, we considered a fixed preconditioner of the form $K_\tau \stackrel{\text{def}}{=} K + \tau M$ for a fixed τ . Suppose we used a different preconditioner at each iteration of the form $K + \tau_k M$ for $k = 1, \dots, m$, then instead of (4) we have,

$$(K + \tau_k M)z_k = v_k \quad k = 1, \dots, m \quad (6)$$

The algorithm is summarized in 4. In this algorithm, in addition to saving V_m , we also save the matrix Z_m . If at every step in the flexible Arnoldi algorithm we use the same value of τ , we are in the same position as in algorithm 1. We have $Z_m = [z_1, \dots, z_m]$, $\bar{H} = \{h_{ik}\}_{1 \leq i \leq m+1, 1 \leq k \leq m}$ and $V_m = [v_1, \dots, v_m]$ which satisfies $V_m^* V_m = I_m$. In addition, we also have the following relations

$$M Z_m = V_{m+1} \bar{H}_m \quad (7)$$

$$K Z_m + M Z_m T_m = V_m \quad (8)$$

where, $T_m = \text{diag}\{\tau_1, \dots, \tau_m\}$. Multiplying (7) by $\sigma_j I_m - T_m$ and adding (8), we obtain for $j = 1, \dots, n_f$

$$(K + \sigma_j M) Z_m = V_{m+1} \underbrace{\left(\begin{bmatrix} I \\ 0 \end{bmatrix} + \bar{H}_m (\sigma_j I_m - T_m) \right)}_{\stackrel{\text{def}}{=} \bar{H}(\sigma_j; T_m)} = V_{m+1} \bar{H}_m(\sigma_j; T_m) \quad (9)$$

We are now in a position to derive a FOM/GMRES algorithm for shifted systems with flexible preconditioning. We search for solutions which are approximations of the form $x_m(\sigma_j) = Z_m y_m(\sigma_j)$, which spans the columns of Z_m . Strictly speaking, $\text{span}\{Z_m\}$ is no longer a Krylov subspace. By minimizing the residual norm over all possible vectors in $\text{span}\{Z_m\}$, we obtain the flexible generalized minimum residual (FGMRES) method for shifted systems, whereas by imposing the Petrov-Galerkin condition $r_m \perp \text{span}\{V_m\}$,

Algorithm 3 Flexible Arnoldi using Modified Gram-Schmidt:

- 1: Given M and b the right hand side, $\tau_j, j = 1, \dots, m$, $v_1 = b/\beta$ and $\beta \stackrel{\text{def}}{=} \|b\|_2$
 - 2: Define the $m+1 \times m$ matrix $\bar{H}_m = \{h_{i,k}\}_{1 \leq i \leq m+1, 1 \leq k \leq m}$.
 - 3: **for all** $k = 1, \dots, m$ **do**
 - 4: Solve $(K + \tau_k M)z_k = v_k$
 - 5: $w_k := Mz_k$
 - 6: **for all** $i = 1, \dots, k$ **do**
 - 7: $h_{ik} := w_k^* v_i$
 - 8: Compute $w_k := w_k - h_{ik} v_k$
 - 9: **end for**
 - 10: $h_{k+1,k} := \|w_k\|_2$. If $h_{k+1,k} = 0$ stop
 - 11: $v_{k+1} = w_k / h_{k+1,k}$
 - 12: **end for**
 - 13: Define $Z_m = [z_1, \dots, z_m]$ and $V_{m+1} = [v_1, \dots, v_{m+1}]$
-

we obtain the flexible full orthogonalized method (FFOM) for shifted systems. This is summarized in algorithm 4. The residuals can be computed as

$$\begin{aligned} r_m(\sigma_j) &= b - (K + \sigma_j M)x_m(\sigma_j) \\ &= V_{m+1}(\beta e_1 - \bar{H}_m(\sigma_k; T_m)y_m(\sigma_j)) \end{aligned} \quad (10)$$

2.2.1 Selecting values of τ_k

In equation (3), at each iteration we solve a system of the form $(K + \tau_k M)z_k = v_k$ for $k = 1, \dots, m$. This cost can be high if the dimension of the Arnoldi subspace m is large and a different preconditioner $K + \tau_k M$ is used at every iteration. In practice, it is not necessary to form and factorize m systems corresponding to different τ_k . In applications, we only need choose a few different τ_k that cover the entire range of the parameters σ_j . The system of equations (6) is solved using a direct solver and since only a few values of τ_k are chosen, the systems can be formed and factorized. Thus, the computational cost will not be affected greatly even if the number of frequencies n_f is large.

Let $\bar{\tau} = \{\bar{\tau}_1, \dots, \bar{\tau}_{n_p}\}$ be the set of values that τ_k can take. In other words, we take n_p distinct preconditioners. Then, the first m_1 values of τ_k are assigned $\bar{\tau}_1$, the next m_2 values of τ_k are assigned $\bar{\tau}_2$ and so on. We also have $m = \sum_{k=1}^{n_p} m_k$.

Algorithm 4 Flexible FOM/GMRES for Shifted Systems

- 1: Given matrices K and M , a right hand side b , $\sigma \in \{\sigma_1, \dots, \sigma_{n_f}\}$
- 2: Choose $T_m = \text{diag}\{\tau_1, \dots, \tau_m\}$ and factorize $K + \tau_k M$.
- 3: **for all** $m = 1, 2, \dots$ **do**
- 4: Run m steps of algorithm 3 to generate V_{m+1} , \bar{H}_{m+1} and Z_m .
- 5: **for all** $k = 1, \dots, n_f$ **do**
- 6: If, not converged
- 7: Construct $\bar{H}_m(\sigma_k; T_m)$ as defined in equation ((9)).
- 8: FOM:

$$H_m(\sigma_k; T_m)y_m^{fom} = \beta e_1$$

- 9: GMRES:

$$y_m^{gm} \stackrel{\text{def}}{=} \min_{y_m \in \mathbb{C}^m} \|\beta e_1 - \bar{H}_m(\sigma_k; T_m)y_m\|_2$$

- 10: Construct the approximate solution as $x_m(\sigma_k) = Z_m y_m$
 - 11: **end for**
 - 12: **end for**
-

2.2.2 Restarting

As the dimension of the subspace m increases, the computational and memory costs increase significantly. A well known solution to this problem is restarting. The old basis is discarded and the Arnoldi algorithm is restarted on a new residual. However, for shifted systems, in order to preserve the shift-invariant property, one needs to ensure collinearity of the residuals of the shifted systems. For FOM the residuals are naturally collinear and the Arnoldi algorithm can be restarted by scaling each residual by some scalar that depends on the shift [26, 12]. For GMRES, the approach used by [10] was extended to shifted systems with multiple preconditioners by [12]. The reader is referred to [12] for further details.

3 Generalized eigenvalue problem and error estimates

We start by computing the approximate eigenvalues and eigenvectors for the matrix $(K + \sigma M)M^{-1}$. Using estimates for approximate eigenvalues and eigenvectors we derive expressions for the convergence of the flexible algorithms. For convenience, we drop the subscript on the shifted frequency,

i.e., use σ instead of σ_j where, $j = 1, \dots, n_f$.

Proposition 1. *Let Z_m , \bar{H}_m and V_{m+1} be computed according to algorithm 3. The eigenpairs of the generalized eigenvalue problem*

$$H_m(\sigma; T_m)z = \theta H_m z \quad (11)$$

then, $\theta, u \stackrel{\text{def}}{=} V_{m+1} \bar{H}_m z$ satisfy the Petrov-Galerkin condition [22, section 4.3.3]

$$(K + \sigma M)M^{-1}u - \theta u \perp \text{span}\{V_m\} \quad u \in \text{span}\{V_{m+1} \bar{H}_m\} \quad (12)$$

Proof. We first begin by manipulating equations (7) and (8). Eliminating Z_m from those equations and adding $\sigma V_{m+1} \bar{H}_m$ to both sides, we have

$$V_{m+1} \bar{H}_m(\sigma_j; T_m) = (K + \sigma M)M^{-1}V_{m+1} \bar{H}_m \quad (13)$$

Now, consider the residual of the eigenvalue calculation for the k th eigenpair, where $k = 1, \dots, m$ is

$$\begin{aligned} r^{\text{eig}}(\sigma) &= (K + \sigma M)M^{-1}V_{m+1} \bar{H}_m z_k - \theta_k V_{m+1} \bar{H}_m z_k \\ &= V_{m+1} \bar{H}_m(\sigma; T_m)z_k - \theta_k V_{m+1} \bar{H}_m z_k \\ &= V_m (H_m(\sigma; T_m)z_k - \theta_k H_m z_k) - h_{m+1,m} v_{m+1}(\tau_m + \theta_k - \sigma) e_m^* z_k \\ &= -h_{m+1,m} v_{m+1}(\tau_m + \theta_k - \sigma) e_m^* z_k \end{aligned} \quad (14)$$

From which we can claim that $u \in \text{span}\{V_{m+1} \bar{H}_m\}$ and $(K + \sigma M)M^{-1}u - \theta u \perp \text{span}\{V_m\}$. In other words, they satisfy the Petrov-Galerkin (12) and are an approximate eigenpair of $(K + \sigma M)M^{-1}$. \square

Furthermore, we define $\rho_k \stackrel{\text{def}}{=} \|(K + \sigma M)M^{-1}u_k - \theta_k u_k\|_2$, which is the residual of the k th eigenvalue calculations. Small ρ_k implies that the eigenvalue calculations have converged. It is readily verified that the eigenvalues λ of KM^{-1} (and the generalized eigenvalue problem $Kx = \lambda Mx$) are related to the eigenvalues $\lambda(\sigma)$ of $(K + \sigma M)M^{-1}$ by the relation $\lambda(\sigma) = \lambda + \sigma$. The importance of the convergence of eigenvalues on the convergence of the Krylov subspace solver using FOM can be established by the following result.

Proposition 2. *The residual using FOM satisfies the following inequality*

$$\|r_m(\sigma)\|_2 \leq \sum_{k=1}^m \rho_k \left| \frac{\sigma - \tau_m}{\theta_k + \tau_m - \sigma} \right| |\theta_k^{-1}| |s_k|$$

where, $s_k \stackrel{\text{def}}{=} e_k^* Z_m^{-1} (H_m^{-1} \beta e_1)$ and ρ_k is the residual of the eigenvalue calculation, defined above.

Proof. We start by writing the residual in equation (10)

$$r_m(\sigma) = V_m(\beta e_1 - H_m(\sigma; T_m)y_m) - v_{m+1}(\sigma - \tau_m)h_{m+1,m}e_m^*y_m$$

Since, for flexible FOM for shifted systems $y_m = H_m(\sigma; T_m)^{-1}\beta e_1$, the first term in the above expression is zero and we have,

$$r(\sigma) = -v_{m+1}(\sigma - \tau_m)h_{m+1,m}e_m^*H_m(\sigma; T_m)^{-1}\beta e_1$$

Now, using the generalized eigendecomposition in (11) $H_m(\sigma; T_m) = H_m Z \Theta Z^{-1}$, and using the residual of the eigenvalue calculation in (14)

$$\begin{aligned} r(\sigma) &= \sum_{k=1}^m v_{m+1}h_{m+1,m}(\sigma - \tau_m)e_m^*z_k\theta_k^{-1}s_k \\ &= \sum_{k=1}^m (r^{\text{eig}}(\sigma)) \frac{\sigma - \tau_m}{\tau_m + \theta_k - \sigma} \theta_k^{-1}s_k \end{aligned} \quad (15)$$

where, $s_k \stackrel{\text{def}}{=} e_k^*Z_k^{-1}H_m^{-1}\beta e_1$. The proof follows from the properties of vector norms. \square

This inequality provides insight into the importance of the accuracy of approximate eigenpairs for the convergence of flexible FOM for shifted systems. We follow the arguments in [19]. In particular, the residual is very small if $\sigma \approx \tau_m$, $|\theta_k^{-1}|$, $|s_k|$ or ρ_k are small. We shall ignore the case that $\sigma \approx \tau_m$ for further analysis, i.e. that the shifted system is exactly the preconditioned system. The eigenvalue residual norm ρ_k being small implies that the Ritz values are a good approximation to the eigenvalues of $(K + \sigma M)M^{-1}$. This implies that all the eigenvalues in this interval have been computed fairly accurately. We now discuss when $|\theta_k^{-1}|$ is large. When all the values of τ_k are equal to τ , the approximate eigenvalues θ_k of KM^{-1} are related to approximate eigenvalues λ_k of the preconditioned system $(K + \sigma M)(K + \tau M)^{-1}$ by the Cayley transformation $\frac{\lambda_k + \sigma}{\lambda_k + \tau}$. Therefore, $|\theta_k^{-1}|$ is large only if $|\lambda_k + \sigma| \ll |\lambda_k + \tau|$. The term s_k can be rewritten as $s_k = e_k^*Z_k^{-1}H_m^{-1}V_m^*V_m\beta e_1 = e_k^*Z_m^{-1}H_m^{-1}V_m^*b$. It is readily verified that $e_k^*Z_m^{-1}H_m^{-1}V_m^*$ is orthonormal to all other approximate eigenvectors $V_{m+1}\bar{H}_mZ_m$ and thus, s_k can be interpreted as the component of the right hand side b in the direction of the approximate eigenvector. In this case, the solution has a small component in the direction of b .

The analysis for the convergence of flexible FOM for shifted systems can be extended to flexible GMRES as well. The following result bounds the difference in the residuals obtained from m steps using flexible FOM and flexible GMRES.

Proposition 3. Let Z_m , \bar{H}_m and V_{m+1} be computed according to algorithm 3. Further, from algorithm 4 we define the flexible FOM quantities $y_m^{\text{fom}}(\sigma) = H_m(\sigma; T_m)^{-1}\beta e_1$, residual $r_m^{\text{fom}}(\sigma) = V_{m+1}(\beta e_1 - \bar{H}_m(\sigma; T_m)y_m^{\text{fom}}(\sigma))$ and flexible GMRES quantities $y_m^{\text{gmres}}(\sigma) = \arg \min_{y \in \mathbb{C}^m} \|\beta e_1 - \bar{H}_m(\sigma; T_m)y\|_2$, residual $r_m^{\text{gmres}}(\sigma) = V_{m+1}(\beta e_1 - \bar{H}_m(\sigma; T_m)y_m^{\text{gmres}}(\sigma))$. Further, assume that $H_m(\sigma; T_m)$ is invertible. We have the following inequality

$$\frac{\|r_m^{\text{fom}}(\sigma) - r_m^{\text{gmres}}(\sigma)\|_2}{\|r_m^{\text{fom}}(\sigma)\|_2} \leq \frac{\|\eta H_m^{-*}(\sigma; T_m)e_m\|_2(1 + \|\eta H_m^{-*}(\sigma; T_m)e_m\|_2)}{1 + \|\eta H_m^{-*}(\sigma; T_m)e_m\|_2^2} \quad (16)$$

where, $\eta \stackrel{\text{def}}{=} h_{m+1,m}(\sigma - \tau_m)$.

Proof. We begin by the following observation from equation (15) $r_m^{\text{fom}}(\sigma) = -\eta e_m^* y_m^{\text{fom}}$ and

$$r_m^{\text{fom}}(\sigma) - r_m^{\text{gmres}}(\sigma) = -V_{m+1}\bar{H}_m(\sigma; T_m)(y_m^{\text{fom}}(\sigma) - y_m^{\text{gmres}})$$

Next, we look at the solution to the GMRES least squares problem which can be written as

$$\bar{H}_m^*(\sigma; T_m)\bar{H}_m(\sigma; T_m)y_m^{\text{gmres}} = \bar{H}_m^*(\sigma; T_m)\beta e_1 = H_m^*(\sigma; T_m)\beta e_1$$

This can be rewritten as

$$\begin{aligned} (H_m^*(\sigma; T_m)H_m(\sigma; T_m) + \eta^2 e_m e_m^*) y_m^{\text{gmres}}(\sigma) &= H_m^*(\sigma; T_m)\beta e_1 \\ (H_m(\sigma; T_m) + \eta^2 H_m^{-*}(\sigma; T_m)e_m e_m^*) y_m^{\text{gmres}}(\sigma) &= \beta e_1 \end{aligned}$$

In other words, the solution to the GMRES subproblem is a rank-one perturbation of the FOM subproblem. Using the Sherman-Morrison identity

$$y_m^{\text{gmres}}(\sigma) = \underbrace{H_m(\sigma; T_m)^{-1}\beta e_1}_{=y_m^{\text{fom}}(\sigma)} - H_m(\sigma; T_m)^{-1}\eta H_m^{-*}(\sigma; T_m)e_m \frac{(\eta e_m^* H_m(\sigma; T_m)^{-1}\beta e_1)}{1 + \|\eta H_m^{-*}(\sigma; T_m)e_m\|_2^2}$$

Then, the residual difference between FOM and GMRES can be bounded as

$$\|r_m^{\text{fom}}(\sigma) - r_m^{\text{gmres}}(\sigma)\|_2 \leq \|\bar{H}_m(\sigma; T_m)H_m(\sigma; T_m)^{-1}\|_2 \|\eta H_m^{-*}(\sigma; T_m)e_m\|_2 \frac{|\eta e_m^* y_m^{\text{fom}}(\sigma)|}{1 + \|\eta H_m^{-*}(\sigma; T_m)e_m\|_2^2}$$

The inequality (16) follows from the following observations $\|\bar{H}_m(\sigma; T_m)H_m(\sigma; T_m)^{-1}\|_2 \leq 1 + \|\eta H_m^{-*}(\sigma; T_m)e_m\|_2$ and $\|r_m^{\text{fom}}\|_2 \leq |\eta e_m^* y_m^{\text{fom}}(\sigma)|$.

□

If $\|\eta H_m^{-*}(\sigma; T_m) e_m\|_2$ is large, then the difference between the two residuals can be large. This happens either when η is large or $H_m(\sigma; T_m)$ is close to singular. In this case, GMRES can stagnate and further progress may not occur. We now discuss situations in which breakdown occurs, i.e. $h_{m+1,m} = 0$. If $h_{m+1,m} \neq 0$ and H_m is full rank, then it can be shown from equation (7) that $\text{span}\{MZ_m\} \subseteq \text{span}\{V_{m+1}\}$ and from equation (9), it follows that $\text{span}\{(K + \sigma_j M)Z_m\} \subseteq \text{span}\{V_{m+1}\}$. Further, $h_{m+1,m} = 0$ if and only if $x_m(\sigma_j)$ is the exact solution and $H_m(\sigma_j; T_m)$ is non-singular. The argument closely follows [23] and will not be repeated here.

4 Inexact preconditioning

We observe that to compute vectors z_k for $k = 1, \dots, m$ in equation (6), we have to invert matrices of the form $K + \tau_k M$. When the problem sizes are large, iterative methods may be necessary to invert such matrices, resulting in a variable preconditioning procedure in which a different preconditioning operator is applied at each iteration. More precisely, for $k = 1, \dots, m$,

$$\tilde{z}_k \approx (K + \tau_k M)^{-1} v_k \quad p_k \stackrel{\text{def}}{=} v_k - (K + \tau_k M) \tilde{z}_k \quad (17)$$

where, p_k is the residual that results after the iterative solver has been terminated. To simplify the discussion, we assume that the termination criteria for the iterative solver is such that $\|p_k\|_2 \leq \varepsilon \underbrace{\|v_k\|_2}_{=1} = \varepsilon$, for some ε . We also

follow closely the approach in [27]. The new flexible Arnoldi relationship is now,

$$(K + \sigma_j M) \tilde{Z}_m + P_m = V_{m+1} \bar{H}_m(\sigma_j; T_m) \quad j = 1, \dots, n_f \quad (18)$$

where, $\tilde{Z}_m = [\tilde{z}_1, \dots, \tilde{z}_m]$ and $P_m = [p_1, \dots, p_m]$ and $\bar{H}_m(\sigma_j; T_m)$ is defined in equation (9). If we seek approximate solutions of the form $x_m(\sigma_j) = \tilde{Z}_m y_m(\sigma_j)$, then we can derive expressions for the approximate residuals as follows,

$$\begin{aligned} r_m(\sigma_j) &= b - (K + \sigma_j M) \tilde{Z}_m y_m(\sigma_j) \\ &= b - V_{m+1} \bar{H}_m(\sigma_j; T_m) y_m(\sigma_j) + P_m y_m(\sigma_j) \\ &= \underbrace{V_{m+1} (\beta e_1 - \bar{H}_m(\sigma_j; T_m) y_m(\sigma_j))}_{\stackrel{\text{def}}{=} \tilde{r}_m(\sigma_j)} + P_m y_m(\sigma_j) \end{aligned}$$

The quantity $r_m(\sigma_j)$ is the true residual and $\tilde{r}_m(\sigma_j)$ is the inexact residual that can be computed by ignoring the error due to early termination of the iterative solver. In practice, computing the true residual (or even the inexact residual) can be expensive. However, in order to monitor the convergence of the iterative solver, we need bounds on the true residual. Using such bounds, we can derive stopping criteria for the flexible Krylov solvers for shifted systems with inexact preconditioning. We start by deriving the difference between the true residual and the inexact residual.

$$\begin{aligned}
\|r_m(\sigma_j) - \tilde{r}_m(\sigma_j)\|_2 &= \|P_m y_m(\sigma_j)\|_2 = \left\| \sum_{k=1}^m e_k^T y_m(\sigma_j) p_k \right\|_2 \\
&\leq \sum_{k=1}^m |e_k^T y_m(\sigma_j)| \|p_k\|_2 \\
&\leq \varepsilon \sum_{k=1}^m |e_k^T y_m(\sigma_j)| = \varepsilon \|y_m(\sigma_j)\|_1
\end{aligned} \tag{19}$$

Using this result, the norm of the true residual $r_m(\sigma_j)$ can be bounded using the following relation

$$\begin{aligned}
\|r_m(\sigma_j)\|_2 &\leq \|r_m(\sigma_j) - \tilde{r}_m(\sigma_j)\|_2 + \|\tilde{r}_m(\sigma_j)\|_2 \\
&\leq \varepsilon \|y_m(\sigma_j)\|_1 + \|\beta e_1 - \bar{H}_m(\sigma_j; T_m) y_m(\sigma_j)\|_2
\end{aligned} \tag{20}$$

Using this bound on the true residual, we can monitor the convergence of iterative solver for each σ_j .

Further, using a similar argument to [27, Proposition 4.1], we can also derive specialized results for the flexible FOM/GMRES for shifted systems. Let $r_m^{\text{fom}}(\sigma_j) \stackrel{\text{def}}{=} b - (K + \sigma_j M) \tilde{Z}_m y_m^{\text{fom}}(\sigma_j)$ and $r_m^{\text{gmres}}(\sigma_j) \stackrel{\text{def}}{=} b - (K + \sigma_j M) \tilde{Z}_m y_m^{\text{gmres}}(\sigma_j)$ be the true residual, respectively resulting from the flexible FOM/GMRES for shifted systems. We have the following error bounds

$$\begin{aligned}
\|V_m^* r_m^{\text{fom}}(\sigma_j)\|_2 &\leq \varepsilon \|y_m^{\text{fom}}(\sigma_j)\|_1 \\
\|(V_{m+1} \bar{H}_m(\sigma_j; T_m))^* r_m^{\text{gmres}}\|_2 &\leq \varepsilon \|\bar{H}_m(\sigma_j; T_m)\|_2 \|y_m^{\text{fom}}(\sigma_j)\|_1
\end{aligned} \tag{21}$$

5 Application to Oscillatory Hydraulic Tomography

In this section, we briefly review the application of Oscillatory Hydraulic Tomography and the Geostatistical approach for solving the resulting inverse problem.

5.1 The Forward Problem

The equations governing ground water flow through an aquifer for a given domain Ω with boundary $\partial\Omega = \partial\Omega_D \cup \partial\Omega_N$, $\partial\Omega_D \cap \partial\Omega_N = \emptyset$ are given by,

$$S_s(\mathbf{x}) \frac{\partial h(\mathbf{x}, t)}{\partial t} - \nabla \cdot (K(\mathbf{x}) \nabla h(\mathbf{x}, t)) = q(\mathbf{x}, t), \quad \mathbf{x} \in \Omega \quad (22)$$

$$h(\mathbf{x}, t) = 0, \quad \mathbf{x} \in \partial\Omega_D \quad (23)$$

$$\nabla h(\mathbf{x}, t) \cdot \mathbf{n} = 0, \quad \mathbf{x} \in \partial\Omega_N \quad (24)$$

where $S_s(\mathbf{x})$ [L^{-1}] represents the specific storage and $K(\mathbf{x})$ [L/T] represents the hydraulic conductivity. In the case of one source oscillating at a fixed frequency ω [radians/T], $q(\mathbf{x}, t)$ is given by

$$q(\mathbf{x}, t) = Q_0 \delta(\mathbf{x} - \mathbf{x}_s) \cos(\omega t) \quad (25)$$

To model periodic simulations, we will assume the source to be a point source oscillating at a known frequency ω and peak amplitude Q_0 at the source location \mathbf{x}_s . In the case of multiple sources oscillating at distinct frequencies, each source is modeled independently with its corresponding frequency as in (25), and then combined to produce the total response of the aquifer.

Since the solution is linear in time, we assume the solution (after some initial time has passed) can be represented as

$$h(\mathbf{x}, t) = \Re(\Phi(\mathbf{x}) \exp(i\omega t)) \quad (26)$$

$\Phi(\mathbf{x})$, known as the phasor, is a function of space only and contains information about the phase and amplitude of the signal. Assuming this solution, the equations (27) in the phasor domain are,

$$\begin{aligned} -\nabla \cdot (K(\mathbf{x}) \nabla \Phi(\mathbf{x})) + i\omega S_s(\mathbf{x}) \Phi(\mathbf{x}) &= Q_0 \delta(\mathbf{x} - \mathbf{x}_s), & \mathbf{x} \in \Omega \\ \Phi(\mathbf{x}) &= 0, & \mathbf{x} \in \partial\Omega_D \\ \nabla \Phi(\mathbf{x}) \cdot \mathbf{n} &= 0, & \mathbf{x} \in \partial\Omega_N \end{aligned} \quad (27)$$

The differential equation (27) along with the boundary conditions are discretized using FEniCS [16, 17, 18] by using standard linear finite elements. Solving it for several frequencies results in system of shifted equations of the form

$$(K + \sigma_j M) x_j = b \quad j = 1, \dots, n_f \quad (28)$$

where, K and M are the stiffness and mass matrices, respectively, that arise precisely from the discretization of (27).

5.2 The Geostatistical Approach

The Geostatistical approach (described in the following papers [13, 15, 14]) is one of the prevalent approaches to solve stochastic inverse problems. The idea is to represent the unknown field as the sum of a few deterministic low-order polynomials and a stochastic term that models small-scale variability. Inference from the measurements is obtained by invoking the Bayes' theorem, through the posterior probability density function which is composed of two parts - likelihood of the measurements and the prior distribution of the parameters. Let $s(\mathbf{x}) \in \mathbb{R}^{N_s}$ be the function to be estimated, here the log conductivity, is modeled by a Gaussian random field. After discretization, the the field can be written as $s \sim \mathcal{N}(X\beta, Q)$. Here X is a matrix of low-order polynomials, β are a set of drift coefficients to be determined and Q is a covariance matrix with entries $Q_{ij} = \kappa(\mathbf{x}_i, \mathbf{x}_j)$, and $\kappa(\cdot, \cdot)$ is a generalized covariance kernel [4]. The measurement equation can be written as,

$$y = h(s) + v, \quad v \sim \mathcal{N}(0, R) \quad (29)$$

where $y \in \mathbb{R}^{N_y}$ represents the noisy measurements and v is a random vector of observation error with mean zero and covariance matrix R . The matrices R , Q and X are part of a modeling choice and more details to choose them can be obtained from the following references [13]. The operator $h : \mathbb{R}^{N_s} \rightarrow \mathbb{R}^{N_y}$ is the known as the parameter-to-observation map or *measurement operator*, with entries h_i that takes the form,

$$h_i = \int_{\Omega} \Re \{ e^{i\omega t} \Phi \delta(\mathbf{x} - \mathbf{x}_i) \} d\Omega \quad (30)$$

where \mathbf{x}_i , is the location of the measurement sensor.

Following the geostatistical method for quasi-linear inversion [13], we compute \hat{s} and $\hat{\beta}$ corresponding to the maximum-a-posteriori probability which is equivalent to computing the solution to a weighted nonlinear least squares problem. To solve the optimization problem, the Gauss-Newton algorithm is used. Starting with an initial estimate for the field s_0 , the procedure is described in algorithm 5.

Algorithm 5 Quasi-linear Geostatistical Approach

1: Compute the $N_y \times N_s$ Jacobian J as,

$$J_k = \left. \frac{\partial h}{\partial s} \right|_{s=s_k} \quad (31)$$

2: Solve the system of equations,

$$\begin{pmatrix} J_k Q J_k^T + R & J_k X \\ (J_k X)^T & 0 \end{pmatrix} \begin{pmatrix} \xi_{k+1} \\ \beta_{k+1} \end{pmatrix} = \begin{pmatrix} y - h(s_k) + J_k s_k \\ 0 \end{pmatrix} \quad (32)$$

3: The update s_{k+1} is computed by,

$$s_{k+1} = X \beta_{k+1} + Q J_k^T \xi_{k+1} \quad (33)$$

4: Repeat steps 1 – 3 until the desired tolerance has been reached. (If necessary, add a line search).

Algorithm 5 requires, at each iteration, computation of the matrices $Q J_k^T$ and $J_k Q J_k^T$. Since the prior covariance matrix Q is dense, a straightforward computation of $Q J_k^T$ can be performed in $\mathcal{O}(N_y N_s^2)$. However, for fine grids, i.e., when the number of unknowns N_s is large, storing Q can be expensive in terms of memory and computing $Q J_k^T$ can be computationally expensive. For regular equispaced grids and covariance kernels that are stationary or translation invariant, an FFT based method can be used to reduce the storage costs of the covariance matrix Q to $\mathcal{O}(N_s)$ and cost of matrix-vector product to $\mathcal{O}(N_s \log N_s)$. For irregular grids, the Hierarchical matrix approach can be used to reduce the storage costs and cost of approximate matrix-vector product to $\mathcal{O}(N_s \log N_s)$ for a wide variety of covariance kernels [25]. Thus, in either situation, the cost for computing $Q J_k^T$ can be done in $\mathcal{O}(N_y N_s \log N_s)$ and the cost of computing $J_k Q J_k^T$ is $\mathcal{O}(N_s N_y \log N_s + N_s N_y)$.

5.3 Sensitivity Matrix computation

Computing the Jacobian matrix J_k at each iteration is often an expensive step. Although explicit analytical expressions for the entries are nearly impossible, several approaches exist. One simple approach is to use finite differences, but this approach is expensive because it requires as many $N_s + 1$ runs of the forward problem, i.e. one more than the number of parameters to be estimated. For large problems, on finely discretized grids the number of unknowns can be quite large, so this procedure is not feasible.

To reduce the computational cost associated with calculating the sensitivity matrix we use the adjoint state method (see for example, [29]). This approach is exact and is computationally advantageous when the number of measurements is far smaller than the number of unknowns. For a complete derivation of the adjoint state equations for the Oscillatory hydraulic tomography, refer to [2]. For the type of measurements described in (30), the entries of the sensitivity matrix can be calculated by the following expression

$$\frac{\partial h_i}{\partial s_j} = \int_{\Omega} \Re \left\{ e^{i\omega t} \left(\left[i\omega \frac{\partial S_s}{\partial s_j} \Phi - \frac{\partial Q_0}{\partial s_j} \right] \Psi + \frac{\partial K}{\partial s_j} \nabla \Phi \cdot \nabla \Psi \right) \right\} d\Omega \quad (34)$$

where Ψ_i , is the known as the *adjoint solution*. It satisfies the following system of equations for each measurement

$$\begin{aligned} -\nabla \cdot (K \nabla \Psi_i) + i\omega S_s \Psi_i &= -\delta(\mathbf{x} - \mathbf{x}_i), & \mathbf{x} \in \Omega \\ \Psi_i &= 0, & \mathbf{x} \in \partial\Omega_D \\ \nabla \Psi_i(\mathbf{x}) \cdot \mathbf{n} &= 0, & \mathbf{x} \in \partial\Omega_N \end{aligned} \quad (35)$$

for $i = 1, \dots, N_y$. The procedure for calculating the sensitivity matrix can thus be summarized as follows. Since (34) is evaluated for all s_j for each

Algorithm 6 Computing Sensitivity Matrix

1. For a given field s_k , solve the forward problem for Φ .
 2. For each measurement i , solve the forward problem for Ψ_i .
 3. Compute the integral in (34) to calculate the sensitivity.
-

measurement i , the adjoint state method requires only $N_y + 1$ forward model solves to compute the sensitivity matrix. Thus, when the number of measurements is far fewer than the number of unknowns, the adjoint state method provides a much cheaper alternative for computing the entries of the Jacobian matrix. This is typically the case in Hydraulic tomography, where having several measurement locations is infeasible because it requires digging new wells.

Further, we realize that equation (34) takes the same form as equation (28) for multiple frequencies. Thus, we can use the algorithms developed in section 2 to solve the system of equations (34) for as many right hand sides as measurements. It is possible to devise algorithms for multiple right hand sides in the context of shifted systems [20, 5] but we will not adopt this approach.

6 Numerical Experiments and Results

We present numerical results for the Krylov subspace solvers and its application to OHT. As mentioned before, we use the FEniCS software [16, 17, 18] to discretize the appropriate partial differential equations. We use the Python interface to FEniCS, with uBLASSparse as the linear algebra back-end. For the direct solvers we use SuperLU [7] package that is interfaced by Scipy whereas for the iterative solver we use an algebraic multigrid package PyAMG [1], with smoothed aggregation along with BiCGSTAB iterative solver. In the following sections, for brevity, we only present results for FOM solver but we observed similar results for the GMRES method as well.

6.1 Krylov subspace solver

In this section, we present some of the results of the algorithms that we have described in section 2. We now describe the test problem that we shall use for the rest of the section. We consider a 2D aquifer in a rectangular domain with Dirichlet boundary conditions on the boundaries. For the log-conductivity field $\log K(\mathbf{x})$, we consider a random field generated using an exponential covariance kernel $\kappa(\mathbf{x}, \mathbf{y}) = 4 \exp(-2\|\mathbf{x} - \mathbf{y}\|_2/L)$. Other parameters used for the model problem are summarized in table 1. We choose 200 frequencies evenly spaced between the minimum and maximum frequencies, which results in 200 systems each of size 90,601.

Definition	Parameters	Values
Aquifer length	L (m)	500
Specific storage	$\log S_s$ (m^{-1})	-11.52
Mean conductivity	$\mu(\log K)$ (m/s)	-11.02
Variance of conductivity	$\sigma^2(\log K)$	1.42
Frequency range	ω (s^{-1})	$[\frac{2\pi}{600}, \frac{2\pi}{3}]$

Table 1: Parameters Chosen For Test Problem

First, we motivate the need for multiple preconditioners to solve the shifted system of equations (28). We begin by looking at the number of iterations taken by restarted GMRES without a preconditioner and using a single preconditioner. We choose a preconditioner of the form $K + \tau M$, for three different values of τ . For illustration purposes, we use a direct solver to invert the preconditioned systems. These values represent the minimum frequency, the average frequency and the maximum frequency in the parameter range. The number of iterations corresponding to restarted GMRES (30) for each of the system is computed and displayed in figure 1. We observe

that the number of iterations corresponding to the systems increases as the frequency of the system decreases. When we use a preconditioner $K + \tau M$, the systems with frequencies nearby $|\tau|$ converge rapidly. However, systems with frequencies further from $|\tau|$ converge more slowly, the further away they are from the frequency of the preconditioned system $|\tau|$. This is consistent with the analysis in section 3 and in particular, proposition 2. Thus, no single preconditioner effectively preconditions all the systems in the given frequency range although, not surprisingly, choosing $|\tau|$ in the center of the frequency range seems to be the best choice. Thus, in order to make the iterative method competitive, we consider using multiple preconditioners to solve the shifted system (28).

We choose preconditioners of the form $K + \tau_k M, k = 1, \dots, m$, where m is the maximum dimension of the Arnoldi iteration. From figure 1, it is clear that the systems with smaller frequencies converge slower, so we choose the values of $|\tau_k|$ that are distributed closer to the origin. In particular, we choose the values of τ_k such that they are evenly spaced on a log scale in the domain $\omega \in [\frac{2\pi}{600}, \frac{2\pi}{3}]$. For example, for $n_p = 5$, the distribution of $|\tau|$ is illustrated in figure 2. Now, let the possible values that $|\tau|$ can take be labeled as $\bar{\tau} = \{\bar{\tau}_1, \dots, \bar{\tau}_{n_p}\}$ where, n_p is the number of distinct number of preconditioner frequencies. Then, the first m_1 values of τ_k are assigned $\bar{\tau}_1$, the next m_2 values of τ_k are assigned $\bar{\tau}_2$ and so on. We also have $m = \sum_{k=1}^{n_p} m_k$. We pick $m_k = m/n_p$. If the algorithm has not converged in m iterations, we restart using the method in section 2.2.2 if we are using a direct solver as the preconditioner. Else, we recycle the same sequence of preconditioners. We implemented both algorithms to invert the preconditioner matrices - using direct solver and using an iterative solver which is an algebraic multigrid preconditioned BiCGSTAB. Using $n_p = 5$ along with the scheme to choose the preconditioner frequencies described above, we observed that the number of iterations (and hence, matrix-vector products) in both cases were less than 40.

Finally, we present the comparison in terms of the run time of our algorithm compared to solving each system using a direct solver. The results are presented in figure 3. In the plots, 'Direct' implies that every system is solved individually by a direct solver. 'Flexible' algorithm uses 5 different preconditioners (see figure 2 with a direct solver for inverting the preconditioners, and solves the FOM subproblem, whereas 'Inexact' uses the preconditioners and inverts the preconditioners using an iterative solver. We see that both the 'Flexible' and 'Inexact' algorithms outperform the 'Direct' approach even for a small number of frequencies. A relative tolerance of $\|r_m(\sigma_j)\|_2 / \|r_0(\sigma_j)\|_2 \leq 10^{-10}$ was used as stopping criteria for all systems $j = 1, \dots, n_f$ and all systems converged within 40 iterations.

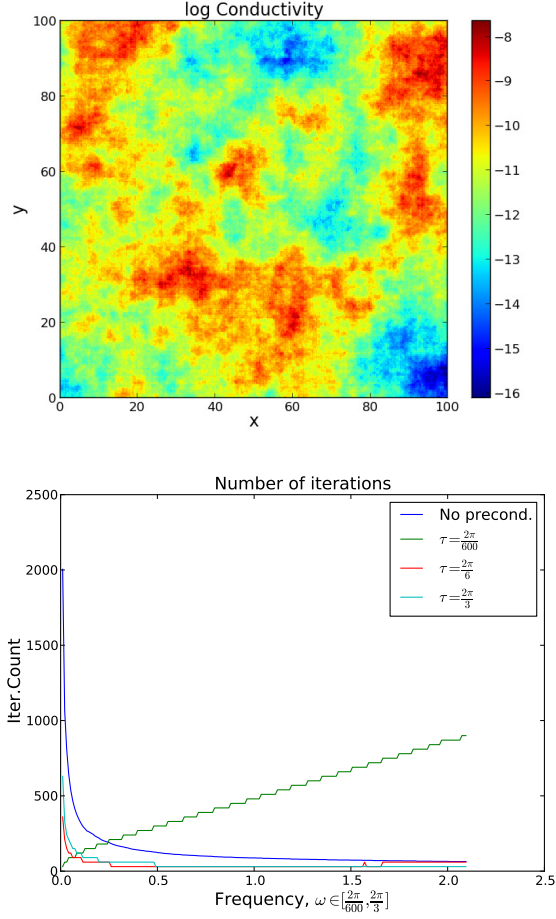


Figure 1: (left) The log conductivity field that we use for the test problem with 90,601 grid points, and (right) iteration count for restarted GMRES (30) for unpreconditioned case and also with single preconditioners of the form $K + \tau M$, with $|\tau| \in \{\frac{2\pi}{600}, \frac{2\pi}{6}, \frac{2\pi}{3}\}$. Results indicate that for the particular choices made, $|\tau| = \frac{2\pi}{6}$ which is roughly at the center of the frequency range, performs best.




Figure 2: Choice of frequencies $|\tau|$ for the preconditioners $K + \tau_k M$. Here, for illustration, we choose $n_p = 5$ and the distinct values of the preconditioned frequencies are evenly space on a log scale in the domain $\omega \in [\frac{2\pi}{600}, \frac{2\pi}{3}]$.

In figures 4 and 5, we plot the residuals and error as a function of the frequency of the system. A direct solver was used for figure 4, whereas an iterative solver (algebraic multigrid preconditioned BiCGSTAB) was used in figure 5 with a stopping critereon $\varepsilon = 10^{-12}$ (refer to section 4. A relative tolerance of $\|r_m(\sigma_j)\|_2 / \|r_0(\sigma_j)\|_2 \leq 10^{-10}$ was used for all systems $j = 1, \dots, n_f$ and all systems converged within 40 iterations. The behavior of residual and the error is quite similar but this is to be expected.

As discussed in section 4, when an inexact preconditioner is used, the flexible Arnoldi relation is no longer exact and the true residual $r_m(\sigma_j)$ and the inexact residual $\tilde{r}_m(\sigma_j)$ are no longer exactly equal. In fact, the error between them can be bounded by the relation (19). In figure, we compare the difference between the true residual and the inexact residuals with the predicted bound $\varepsilon \|y_m(\sigma_j)\|_1$. We see 6 that the bound is fairly accurate. The stopping tolerances were chosen to be $\varepsilon = 10^{-10}, 10^{-12}, 10^{-14}$. In fact, for larger stopping tolerances for the inner solver, the outer solver did not converge.




Figure 3: Comparison of time for the different algorithms. 'Direct' implies that every system is solved individually by a direct solver. 'Flexible' algorithm uses 5 different preconditioners (see figure 2 with a direct solver for inverting the preconditioners, and solves the FOM subproblem, whereas 'Inexact' uses the preconditioners and inverts the preconditioners using an iterative solver. We see that both the 'Flexible' and 'Inexact' algorithms outperform the 'Direct' approach even for a small number of frequencies.




Figure 4: (left) residual vs frequencies for FOM and GMRES and (right) error vs frequency for FOM and GMRES. All systems converged within 40 iterations. A tolerance of $\|r_m(\sigma_j)\|_2 / \|r_0(\sigma_j)\|_2 \leq 10^{-10}$ was used for all systems $j = 1, \dots, n_f$. The preconditioner matrices were inverted using a direct solver.




Figure 5: (left) residual vs frequencies for FOM and GMRES and (right) error vs frequency for FOM and GMRES. All systems converged within 40 iterations. A tolerance of $\|r_m(\sigma_j)\|_2/\|r_0(\sigma_j)\|_2 \leq 10^{-10}$ was used as stopping criteria for all systems $j = 1, \dots, n_f$. The preconditioner matrices were inverted using an iterative solver with $\varepsilon = 10^{-12}$ as an inner stopping criterion.




Figure 6: Difference in the norm between the true and the inexact residuals, when an iterative solver is used as the preconditioner. Three different stopping tolerances were considered $\varepsilon = 10^{-10}, 10^{-12}, 10^{-14}$.

6.2 Application: Tomographic reconstruction

The objective is now to determine the random conductivity field $K(\mathbf{x})$ from discrete measurements of the head h obtained from several pumping tests performed with multiple frequencies. Since the conductivity field needs to be positive, so that the forward problem is well-posed, we consider a log-transformation $s = \log K$. The "true" field is taken to be that in figure 7 which is a scaled version of Franke's function [8]. We choose the covariance matrix Q to have entries $Q_{ij} = \kappa(\mathbf{x}_i, \mathbf{x}_j)$, corresponding to an exponential covariance kernel

$$\kappa(\mathbf{x}, \mathbf{y}) = \exp\left(-\frac{\|\mathbf{x} - \mathbf{y}\|_2}{4L}\right)$$

where, L is the length of the domain. We also choose $R = \sigma^2 I$ and $X = [1, \dots, 1]^T$.

The measurements are obtained by taking as the true log conductivity field, the field in figure 7. Then, the phasor is calculated by solving equations (27) with the source location and measurement locations given in 7. The measurements are collected for each frequency and two pieces of information are recored, the coefficients of the sine and cosine terms in equation (26). The inverse problem is then solved using these measurements. The frequency range that is chosen is $\omega \in [2\pi/150, 2\pi/30]$.

The time for computing the Jacobian is listed in figure 8. For the iterative solver, we use the flexible FOM solver using direct solver as preconditioner. The relative stopping tolerance we used was 10^{-10} . Since the cost to solve systems with multiple frequencies is nearly the same as the cost to solve a single system, the time for building the Jacobian is, more or less, independent of the number of frequencies. However, when a direct solver is used to independently solve the systems for multiple frequencies, the cost for constructing the Jacobian scales linearly with the number of frequencies. This results in significant reduction in the cost for solving the inverse problem, since constructing the Jacobian is the most expensive part of solving the inverse problem. The disparity in the computation times for the Jacobian between the direct approach and the iterative procedure is exacerbated further, with larger problem sizes resulting from finer discretizations.

Finally, we compare the error in the the reconstruction with multiple frequencies. Table 2 lists the L^2 error in the reconstruction. We report two errors - the first being the L^2 error in the entire domain, the second being the L^2 error in the area enclosed by the measurement wells. While increasing the number of frequencies does not seem to improve the error in the entire domain, it improves the error in the region enclosed by the measurement wells. This is the primary motivation for using multiple frequencies in the




Figure 7: (left) the true log conductivity field that we use for the synthetic inverse problem, and (right) location of measurement wells and the source.

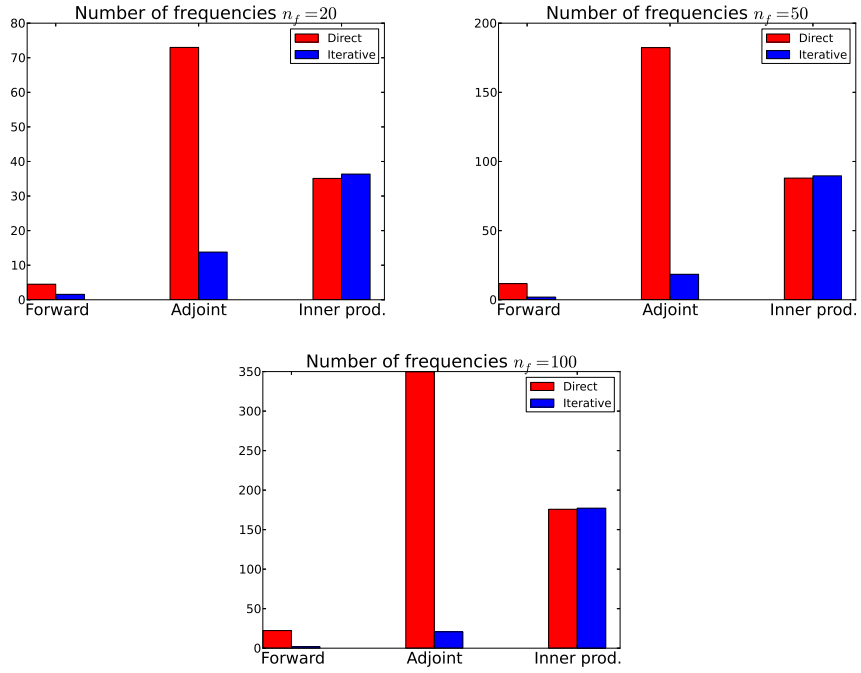


Figure 8: Comparison of time taken for different components in the Jacobian. "Forward" refers to solving the forward problem for multiple frequencies, equation (27). "Adjoint" refers to solving the adjoint field for multiple frequency at each measurement location, equation (35). "Derivatives" refer to forming the inner product to form the rows of the Jacobian, equation (34).

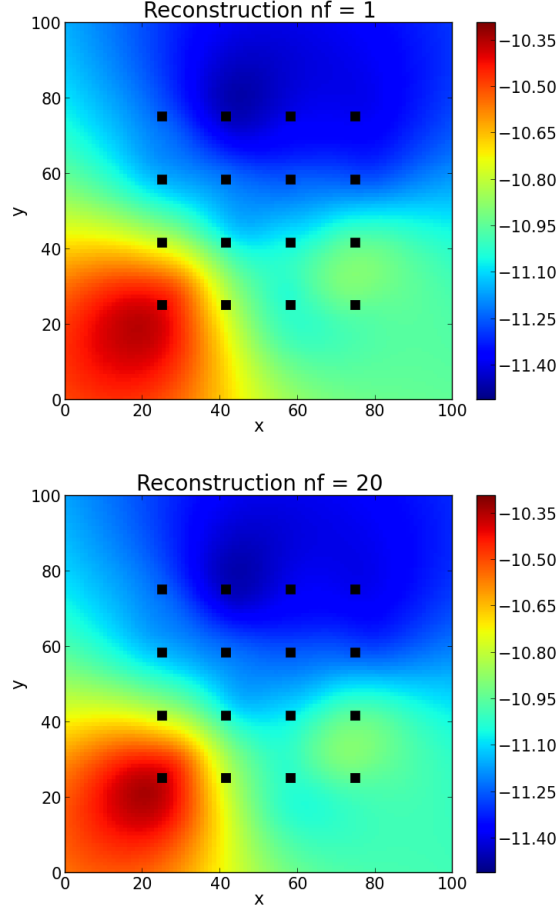


Figure 9: Comparison of inversion results for log conductivity with $n_f = 1$ and $n_f = 20$ frequencies. The errors in reconstruction are reported in 2.

inversion. However, beyond a point, the addition of frequencies does not seem to reduce the error. This might be because there is no additional information that is obtained from the addition of measurements with these frequencies, and, to further improve estimation accuracy, one would need to introduce more stimulation and observation points [2].

7 Conclusions

We have presented a flexible Krylov subspace algorithm for shifted systems of the form (28) that uses multiple shifted preconditioners of the form $K + \tau M$. The values of τ are chosen in order to improve convergence of the solver for

N_f	Total error	Error within box
1	0.2604	0.0622
5	0.2387	0.0554
10	0.2263	0.0509
20	0.2170	0.0461

Table 2: L^2 error due to the reconstruction. We report two errors - the first being the L^2 error in the entire domain, the second being the L^2 error in the area enclosed by the measurement wells.

all the shifted systems. The number of preconditioners chosen varies based on the distribution of the shifts. A good rule of thumb is that the systems having shift σ will converge faster if there a preconditioner with shift τ that is nearby σ . When the size of the linear systems is much larger, direct solvers are much more expensive. In such cases, preconditioning would be done using iterative solvers. The error analysis in section 4 provides insight into monitor approximate residuals without constructing the true residuals. One can naturally extend the ideas in this paper to systems with multiple shifts and multiple right hand sides using either block or deflation techniques.

We applied the flexible Krylov solver to an application problem that benefited significantly from fast solvers for shifted systems. In particular, oscillatory hydraulic tomography is an emerging technique for aquifer characterization. However, since drilling observation wells to obtain measurements is expensive, oscillatory hydraulic tomography aims to generate more informative measurements by pumping at different frequencies using the same pumping locations and measurement wells. In future studies we aim to study more realistic conditions for tomography, including a joint inversion for storage and conductivity. This would be ultimately beneficial to the practitioners. We envision that fast solvers for shifted systems would be beneficial for rapid aquifer characterization.

References

- [1] W. N. Bell, L. N. Olson, and J. B. Schroder. PyAMG: Algebraic multi-grid solvers in Python v2.0, 2011. Release 2.0.
- [2] M. Cardiff, T. Bakhos, P.K. Kitanidis, and W. Barrash. Multi-frequency oscillatory hydraulic tomography: The use of steady-state periodic signals for sensitivity analysis and inversion. *Water Resources Research*, in review.

- [3] M. Cardiff and W. Barrash. 3-d transient hydraulic tomography in unconfined aquifers with fast drainage response. *Water Resources Research*, 47(12):W12518, 2011.
- [4] G. Christakos. On the problem of permissible covariance and variogram models. *Water Resources Research*, 20(2):251–265, 1984.
- [5] D. Darnell, R.B. Morgan, and W. Wilcox. Deflated gmres for systems with multiple shifts and multiple right-hand sides. *Linear Algebra and its Applications*, 429(10):2415–2434, 2008.
- [6] B.N. Datta and Y. Saad. Arnoldi methods for large sylvester-like observer matrix equations, and an associated algorithm for partial spectrum assignment. *Linear Algebra and its Applications*, 154:225–244, 1991.
- [7] James W. Demmel, Stanley C. Eisenstat, John R. Gilbert, Xiaoye S. Li, and Joseph W. H. Liu. A supernodal approach to sparse partial pivoting. *SIAM J. Matrix Analysis and Applications*, 20(3):720–755, 1999.
- [8] R. Franke. A critical comparison of some methods for interpolation of scattered data. Technical report, DTIC Document, 1979.
- [9] R.W. Freund. Solution of shifted linear systems by quasi-minimal residual iterations. *Numerical Linear Algebra*, pages 101–121, 1993.
- [10] A. Frommer and U. Glässner. Restarted gmres for shifted linear systems. *SIAM Journal on Scientific Computing*, 19(1):15–26, 1998.
- [11] G.H. Golub and C.F. Van Loan. *Matrix computations*, volume 3. Johns Hopkins University Press, 1996.
- [12] G. Gu, X. Zhou, and L. Lin. A flexible preconditioned arnoldi method for shifted linear systems. *Journal of Computational Mathematics - International Edition*, 25(5):522, 2007.
- [13] P. K. Kitanidis. Quasilinear geostatistical theory for inversing. *Water Resour. Res.*, 31(10):2411–2419, 1995.
- [14] P. K. Kitanidis. *On stochastic inverse modeling*, volume Geophysical Monograph 171, pages 19–30. AGU, Washington, D. C., 2007.
- [15] P. K. Kitanidis. *Bayesian and Geostatistical Approaches to Inverse Problems*, pages 71–85. John Wiley & Sons, Ltd, 2010.

- [16] Anders Logg, Kent-Andre Mardal, Garth N. Wells, et al. *Automated Solution of Differential Equations by the Finite Element Method*. Springer, 2012.
- [17] Anders Logg and Garth N. Wells. Dofin: Automated finite element computing. *ACM Transactions on Mathematical Software*, 37(2), 2010.
- [18] Anders Logg, Garth N. Wells, and Johan Hake. *DOLFIN: a C++/Python Finite Element Library*, chapter 10. Springer, 2012.
- [19] K. Meerbergen. The solution of parametrized symmetric linear systems. *SIAM Journal on Matrix Analysis and Applications*, 24(4):1038–1059, 2003.
- [20] K. Meerbergen and Z. Bai. The lanczos method for parameterized symmetric linear systems with multiple right-hand sides. *SIAM Journal on Matrix Analysis and Applications*, 31(4):1642–1662, 2010.
- [21] M. Papolizio and V. Simoncini. Acceleration techniques for approximating the matrix exponential. *SIAM J. Matrix Analysis Appl.* v30 i2, pages 657–683, 2008.
- [22] Y. Saad. *Numerical methods for large eigenvalue problems*, volume 158. SIAM, 1992.
- [23] Y. Saad. A flexible inner-outer preconditioned gmres algorithm. *SIAM Journal on Scientific Computing*, 14:461–461, 1993.
- [24] Y. Saad. *Iterative methods for sparse linear systems*. Society for Industrial Mathematics, 2003.
- [25] A.K. Saibaba and P.K. Kitanidis. Efficient methods for large-scale linear inversion using a geostatistical approach. *Water Resources Research*, 48(5):W05522, 2012.
- [26] V. Simoncini. Restarted full orthogonalization method for shifted linear systems. *BIT Numerical Mathematics*, 43(2):459–466, 2003.
- [27] V. Simoncini and D.B. Szyld. Theory of inexact krylov subspace methods and applications to scientific computing. *SIAM Journal on Scientific Computing*, 25(2):454–477, 2003.
- [28] V. Simoncini and D.B. Szyld. Recent computational developments in krylov subspace methods for linear systems. *Numerical Linear Algebra with Applications*, 14(1):1–59, 2007.

- [29] N.Z. Sun, W.W.G. Yeh, et al. Coupled inverse problems in groundwater modeling 1. sensitivity analysis and parameter identification. *Water Resources Research*, 26(10):2507–2525, 1990.